

Symmetry-Protected Local Minima in Infinite DMRG

Robert N. C. Pfeifer*

Dept. of Physics & Astronomy, Macquarie University, Sydney, NSW 2109, Australia

(Dated: August 19, 2015)

The infinite Density Matrix Renormalisation Group (iDMRG) algorithm is a highly successful numerical algorithm for the study of low-dimensional quantum systems, and is also frequently used to initialise the more popular finite DMRG algorithm. Implementations of both finite and infinite DMRG frequently incorporate support for the protection and exploitation of symmetries of the Hamiltonian. In common with other variational tensor network algorithms, convergence of iDMRG to the ground state is not guaranteed, with the risk that the algorithm may become stuck in a local minimum. In this paper I demonstrate the existence of a particularly harmful class of physically irrelevant local minima affecting both iDMRG and to a lesser extent also infinite Time-Evolving Block Decimation (iTEBD), for which the ground state is compatible with the protected symmetries of the Hamiltonian but cannot be reached using the conventional iDMRG or iTEBD algorithms. I describe a modified iDMRG algorithm which evades these local minima, and which also admits a natural interpretation on topologically ordered systems with a boundary.

I. INTRODUCTION

The Density Matrix Renormalisation Group (DMRG) algorithm¹⁻³ is a numerical algorithm for computing the ground state of a lattice Hamiltonian on a finite or infinite chain, and is one of the most successful and widespread algorithms in condensed matter physics.⁴ In the last two decades it has also been recognised that both finite and infinite DMRG may be understood as tensor network algorithms, with the ground state being expressed in the form of a Matrix Product State (MPS).⁴⁻⁷ Implementations of DMRG typically exploit superselection principles with respect to quantum numbers such as particle number or total spin in order to protect these symmetries and/or increase computational efficiency,⁷⁻⁹ and in the tensor network formulation of the DMRG algorithm this equates to working with tensors which explicitly respect and exhibit the corresponding global symmetries of the system.¹⁰⁻¹³

It has been known for some time now that in the infinite form of the DMRG algorithm (iDMRG), complications may arise in systems having non-bosonic statistics. Such systems typically converge to a stable orbit in the space of MPS states rather than to a single state, with the spread of the energies of these states being determined by the truncation accuracy of the MPS.^{8,14,15} In the present paper I demonstrate a more serious problem which may arise in iDMRG simulations with non-trivial particle statistics, if these statistics are also exploited for computational advantage: When the MPS tensors of the true ground state act as intertwinors on the space of protected symmetry charge labels with a period $p > 2$, the initial boundary conditions of the simulation may force it to converge to a state which is not (and does not closely approximate) the ground state.

In Sec. II of this paper I review the iDMRG algorithm, while in Sec. III I describe how the performance of the infinite DMRG algorithm may suffer when symmetries of the Hamiltonian are exploited for computational advantage, and give a simple example in which the MPS

tensors of the ground state act as period-3 intertwinors. In Sec. III E I show how the iDMRG algorithm may be modified to circumvent this problem and avoid these local minima, while Sec. III F discusses how the improved algorithm has a natural physical interpretation when applied to topologically ordered systems with a boundary.¹⁶⁻¹⁹ A more realistic example is given in Sec. IV, which demonstrates the superior convergence of the revised algorithm. Finally Sec. V argues that equivalent local minima are expected to exist for the iTEBD algorithm, though this algorithm is seen to avoid most symmetry-protected local minima as a result of the boundary conditions implicit when the simulation is initialised.

II. INFINITE DMRG

A. Overview

Behind every tensor network algorithm lies a tensor network Ansatz, and for DMRG, this Ansatz is the MPS. For an infinite system this Ansatz is likewise, in principle, infinite in extent, as per Fig. 1(i). In practice the limit of an infinite Ansatz is approached asymptotically. The iDMRG algorithm is initialised by exactly diagonalising a two-site Hamiltonian and decomposing the lowest-energy eigenvector v into three components, $A^{(1)}$, $\lambda^{(1)}$, and $B^{(1)}$, as shown in Fig. 1(ii)-(iii). Note that $A^{(1)}$ and $B^{(1)}$ are typically gauged to be left-unitary and right-unitary respectively, as shown in Fig. 1(iv), and $\lambda^{(1)}$ is diagonal and non-negative. Tensors $A^{(1)}$, $B^{(1)}$, and $\lambda^{(1)}$ thus correspond to the Schmidt (or singular value) decomposition of the eigenvector v . Iteration 2 of the algorithm then consists of taking the centre two sites of the Ansatz, inserting an additional two sites in between them, and computing a new triplet of tensors $A^{(2)}$, $\lambda^{(2)}$, and $B^{(2)}$ so as to minimise the total energy of the Hamiltonian (now on four sites), subject to fixed $A^{(1)}$ and $B^{(1)}$ and constraints on the maximum permitted index dimensions of the tensors. Tensors $A^{(2)}$, $B^{(2)}$, and $\lambda^{(2)}$ are gauged




FIG. 1. (i) The infinite Matrix Product State, in tensor network and traditional DMRG notations. Orientations are applied to all indices of the MPS to enable the exploitation of quantum numbers which are not self-dual. (ii) Construction of the initial two-site MPS. Object v is the lowest-energy eigenvector of the two-site Hamiltonian with eigenvalue E_0 , and is decomposed into three components, $A^{(1)}$, $\lambda^{(1)}$, and $B^{(1)}$, where the superscript (1) indicates that these are the tensors constructed during iteration 1 of the algorithm. The A and B tensors typically have three indices apiece, so we may also represent this as shown in diagram (iii) where the additional indices are of dimension 1 and carry trivial quantum number (also described as “the vacuum charge”, and frequently denoted \mathbb{I}). (iv) All A and B tensors satisfy left-unitarity and right-unitarity respectively.

as before.

This tensor pair insertion procedure is then iterated repeatedly until a steady state is reached such that $\lambda^{(k)} \approx \lambda^{(k-1)}$ for all $k > k_{\text{convergence}}$, to within some desired tolerance

$$\sum_i \left(\lambda^{(k)}_i - \lambda^{(k-1)}_i \right)^2 \leq \varepsilon, \quad (1)$$

or that the values of the tensors $\lambda^{(k)}$ are cyclic with some integer period p for all $k > k_{\text{convergence}}$, i.e. $\lambda^{(k)} \approx \lambda^{(k-p)}$ for all k after convergence has been attained, again within some tolerance ε . Formally, exact convergence corresponds to $\varepsilon \rightarrow 0$, but in practice this is limited by the numeric precision of the software implementation. This process is illustrated in Fig. 2.

Note that there exists an additional gauge freedom in the choice of $A^{(i)}$, $\lambda^{(i)}$, and $B^{(i)}$ which is not fixed by the above prescription. The decomposition

$$v \rightarrow \{A^{(i)}, \lambda^{(i)}, B^{(i)}\} \quad (2)$$




FIG. 2. One iteration of the infinite DMRG algorithm, (i) in tensor network notation and (ii) in traditional DMRG notation.

is invariant under a transformation of the form

$$\begin{aligned} A^{(i)} &\rightarrow A^{(i)} U \\ B^{(i)} &\rightarrow U^\dagger B^{(i)} \end{aligned} \quad (3)$$

where U is a block-diagonal unitary matrix acting on the index shared with $\lambda^{(i)}$, whose block sizes correspond to the degeneracy of the associated singular values in $\lambda^{(i)}$. Thus, for example, if all entries in $\lambda^{(i)}$ are unique, then U is a diagonal matrix of complex phases. This gauge may be changed at any time by inserting a pair of matrices $U U^\dagger$ on an internal bond of the MPS and absorbing U and U^\dagger into the A or B tensors to the left and right of the bond.

On convergence, there exists a choice of gauge as per Eq. (3) such that an infinite MPS may be constructed by infinitely repeating blocks made up of the last p tensors on each side,

$$\lim_{n \rightarrow \infty} \left[A^{(k-p+1)} \dots A^{(k)} \right]^n \lambda^{(k)} \left[B^{(k)} \dots B^{(k-p+1)} \right]^n, \quad (4)$$

and optionally absorbing matrix $\lambda^{(k)}$ into either $A^{(k)}$ or $B^{(k)}$ to obtain the form of Fig. 1(i).²⁰ The low-energy state to which the iDMRG algorithm has thus converged is then assumed to closely approximate the ground state of the infinite system.

In order for this assumption to be valid, it is necessary (though not sufficient)

1. that the ground state of the system may be well-approximated (or exactly represented) by an infinite MPS, and
2. that the iDMRG algorithm is capable of converging to that MPS.

Where multiple degenerate ground states exist which may be written in MPS form, these are associated with a degeneracy of singular values in $\lambda^{(k)}$ and therefore related by a choice of MPS gauge. In such a case it is therefore always possible to write an orthogonal basis for the ground subspace with each basis vector taking the form of Eq. (4).

B. A note on notation

Over the years, a number of different nomenclatures have been adopted for the tensors of the MPS Ansatz. In this paper, when MPS tensors are gauged according to Fig. 2 (left- or right-unitary gauge) they are denoted A or B respectively. Tensors which are not necessarily in the left-unitary or right-unitary gauge have historically been denoted by either A or Γ ; this paper will use Γ to avoid ambiguity.

Regarding arrows on the indices of the MPS, these arrows have historically been used for two purposes. They may either signify renormalisation group flow, with the A and B tensors performing a coarse-graining on the Hilbert space of the lattice and λ being a matrix representation of a ket acting on a Hilbert space of dimension $D \times D$ (where D is the maximum permitted dimension of the MPS bond indices), or they may indicate an orientation for the summing of charges in simulations with protected symmetries.¹⁰ In the former, it is conventional for all arrows to flow towards λ and thus those on the MPS bond indices of the B tensors would point from right to left. The latter imposes no strict constraint on orientation, requiring only that reversal of an arrow be associated with replacing each charge on that index with its dual.

While the $AB\lambda$ form of the MPS Ansatz [e.g. Fig. 2(i)] is of great practical use when implementing the iDMRG algorithm, writing an MPS in this form may conceal relevant structure as the Ansatz itself explicitly breaks translation invariance. This drawback is avoided when using the Γ form of the Ansatz [e.g. Figs. 1(i) and 3], in which a state which is invariant under translation by p sites may be represented by an MPS made up by periodically repeating p different tensors $\Gamma_1 \dots \Gamma_p$. As this paper makes use of both the $AB\lambda$ and the Γ forms of the Ansatz, a convention is adopted in which all bond indices flow from left to right—both in the $AB\lambda$ form and in the Γ form—for greater consistency between the two notations.

Finally, a wide variety of terminology is used to describe the different regions of the MPS chain when performing DMRG. Historically, in the $AB\lambda$ form of the Ansatz the A portion of the chain, to the left of λ and

excluding any new tensors being inserted or updated, has been termed the *system block* while the B portion (to the right of λ , excluding any new tensors being inserted or updated) has been termed the *environment block*. When the tensors being inserted or updated are included, these regions become the *system superblock* and *environment superblock* respectively. The DMRG algorithm is symmetric with respect to the A and B portions of the chain, and so from an MPS perspective the distinction between the system and the environment is only of historical importance, but the distinction between block and superblock remains useful.

III. LOCAL MINIMA FROM SYMMETRY PROTECTION

A. Illustrative example

As a simple example of where the second assumption of Sec. II A fails, and the iDMRG algorithm is incapable converging to the true ground state of the system, consider the Hamiltonian

$$\hat{H} = - \sum_i \mu \hat{n}_i \quad (5)$$

on a chain populated by hard-core particles, and introduce a conserved quantum number corresponding to particle number modulo 3. This may seem artificial in the context of Eq. (5), but arises naturally when a chemical potential having the form of Eq. (5) is attached to a pre-existing Hamiltonian having \mathbb{Z}_3 symmetry, in the limit that μ becomes large (see e.g. Sec. IV and Ref. 21).

For $\mu > 0$, the Hamiltonian (5) has a unique ground state corresponding to complete filling of the chain. This state admits an MPS representation consistent with protection of the \mathbb{Z}_3 symmetry, as shown in Fig. 3(i), and has bond dimension $D = 1$.

In this representation, if the MPS bond index between sites i and $i + 1$ carries charge 0, then the index between $i + 1$ and $i + 2$ carries charge 1, and more generally, that between $i + q$ and $i + q + 1$ carries charge $(q \bmod 3)$. This MPS is not translation-invariant, but we may make it so by defining \hat{T} to be the one-site translation operator and writing

$$|\psi'\rangle = \frac{1}{\sqrt{3}} \left(|\psi\rangle + \hat{T}|\psi\rangle + \hat{T}^2|\psi\rangle \right), \quad (6)$$

with MPS bond dimension $D = 3$. The factors of $1/\sqrt{3}$ are incorporated by arbitrarily selecting a single MPS bond and introducing an additional tensor on this bond, $\lambda = \text{diag}(1/\sqrt{3}, 1/\sqrt{3}, 1/\sqrt{3})$. While it would be conventional to denote tensors to the left of this bond by $A^{(k)}$, $A^{(k-1)}$, \dots , and tensors to the right by $B^{(k)}$, $B^{(k-1)}$, \dots , all of these tensors are in fact identical. In Fig. 3 they are therefore all denoted Γ to emphasise this fact. The new state is invariant under the operation

$$\lambda \rightarrow \Gamma \lambda \Gamma, \quad (7)$$

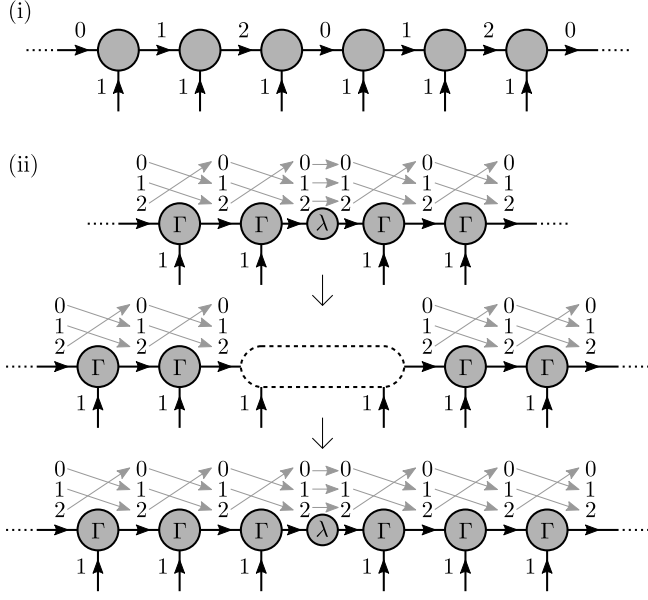


FIG. 3. (i) Ground state of Eq. (5) with three-site translation invariance, requiring MPS bond index dimension $D = 1$. Index labels correspond to charge sectors of non-zero dimension. (ii) Insertion of two extra sites into $|\psi'\rangle$ of Eq. (6). Grey arrows show how the tensor Γ acts as an example on the charge sectors of the MPS bond.

which inserts an additional two sites (but leaves the MPS unchanged, as it is infinite). This is represented graphically in Fig. 3(ii). State $|\psi'\rangle$ therefore in principle represents a possible steady state of the infinite DMRG algorithm. In practice, however, this state cannot be attained by the algorithm in its usual form.

This is readily seen by considering Fig. 1(iii). On this first iteration, the vacuum charge boundary conditions may only be satisfied if the number of particles on the sites is a multiple of three. The hard-core constraint sets an upper bound of two, and thus the Ansatz is forced into a zero-particle state. The bonds between $A^{(1)}$, $\lambda^{(1)}$, and $B^{(1)}$ in Fig. 1(iii) are therefore also in the vacuum charge sector, and this constraint propagates to the next iteration. At each step it therefore remains impossible to break out of the zero-particle sector. On the space of states which can be explored by the iDMRG algorithm with the conventional choice of boundary conditions, this state consequently represents a “local” minimum which is not the ground state of the Hamiltonian (5).²²

Note that if the \mathbb{Z}_3 symmetry is not enforced, then the iDMRG algorithm does not become trapped in this minimum. Enforcing symmetries reduces the space of possible states to which the algorithm may evolve in one iteration. If (as here) this restriction directly results in the algorithm becoming stuck in a state which is not the true ground state, this may be termed a *symmetry-protected local minimum*. In this instance the local minimum is also associated with complete filling of the lattice, but this is not a prerequisite and in Sec. IV we




FIG. 4. (i) Tensor P is a symmetry-conserving projection operator acting on the MPS bond index, reducing the dimension of the associated Hilbert space from D to D' . (ii) Tensor P acts on an MPS tensor Γ . (iii) Calculation of the unnormalised 1-site reduced density matrix associated with P , denoted $\hat{\rho}^{(P)}$.

shall see a more realistic Hamiltonian which also exhibits symmetry-protected local minima at large but incommensurate fillings.

B. MPS tensors as intertwinors

To see how the tensors Γ in Fig. 3 may be understood as intertwinors, note that under Hamiltonian (5) with $\mu > 0$ the probability amplitude of any site of the lattice being unoccupied is zero. We may therefore project the physical indices of these tensors into the dimension-1 subspace corresponding to the presence of a particle, and the MPS tensors then reduce to matrices. Conservation of charge requires that these matrices act as an intertwinor on the charge sectors of the MPS bond index.

More generally, let Γ be a translation-invariant representation of the ground state of a translation-invariant Hamiltonian \hat{H} ,²³ let P be a symmetry-conserving projector acting on the MPS bond index which reduces the dimension of this index from D to some smaller value D' , and let $\mathcal{Q} = \{q_P\}$ be the space of charges appearing on the smaller index of P . Let $\hat{\rho}^{(P)}$ be the unnormalised 1-site reduced density matrix constructed from the projected MPS tensor $P^\dagger \Gamma_i P$, where i is the physical index, as shown in Fig. 4. Suppose there exists a choice of P such that $[\hat{\rho}^{(P)}]_i^i$ (no sum over i) is non-vanishing for only one value of i , denoted i_P , and that value i_P is associated with a charge having non-trivial action on \mathcal{Q} . It then follows that $P^\dagger \Gamma_{i_P} P$ is an intertwinor on the space of charges \mathcal{Q} . If there exists a set of such projectors \mathcal{P} such that $\text{Tr}(PP'^\dagger) = 0$ for all pairs $(P, P') \in \mathcal{P}$, and for which the value of

$$\hat{\rho}^{(\mathcal{P})} = \sum_{P \in \mathcal{P}} [\hat{\rho}^{(P)}]_{i_P}^{i_P} \quad (8)$$

is sufficiently greater than zero, then this indicates that

$$\bigoplus_{P \in \mathcal{P}} [\hat{\rho}^{(P)}]_{i_P}^{i_P} \quad (9)$$

makes a significant contribution to the one-site reduced




FIG. 5. Calculation of the initial value of v with multiple boundary charges. The left index is permitted to take $n_q^{(L)}$ different charges, each of degeneracy 1, and the right index is permitted to take $n_q^{(R)}$ different charges, again each with degeneracy 1. These need not be the same charges as appear on the left. This diagram reduces to Fig. 1(ii)-(iii) for $n_q^{(L)} = n_q^{(R)} = 1$ and $q_L = q_R = \mathbb{I}$.

density matrix, and the object

$$\bigoplus_{P \in \mathcal{P}} P^\dagger \Gamma_{i_P} P, \quad (10)$$

which is an intertwinor on a subspace of the MPS bond index, makes a significant contribution to the construction of the ground state. We may then anticipate the existence of a symmetry-protected local minimum.

Having identified the conditions which may cause the iDMRG algorithm to become stuck in a symmetry-protected local minimum, the following Sections explore different strategies for modifying (i) the initial conditions (Sec. III C), and (ii) the iDMRG update algorithm (Secs. III D and III E), in order to either avoid falling into, or permit escape from, the vacuum sector. In the former approach, the zero-particle state with trivial charges on the MPS bond indices remains a local minimum as once the MPS bond enters this charge configuration, the update algorithm is unable to escape. The latter approach is more robust, eliminating the local minimum from the landscape of states by providing a means for the algorithm to evolve to a lower-energy state, while nevertheless retaining protection of the global \mathbb{Z}_3 symmetry.

C. Infinite DMRG with free boundary charges

One way in which we can avoid the local minima is by allowing multiple choices of boundary charge on the initial state. Object v continues to satisfy $vH = \lambda H$ for the initial 2-site Hamiltonian, but this expression now takes the form shown in Fig. 5. For the example of Sec. III A, allowing both the left and the right boundary charges to take any value from $\{0, 1, 2\}$ results in the iDMRG algorithm finding the translation-invariant ground state $|\psi'\rangle$.

It is interesting to compare this approach with the technique used in finite DMRG to compute the lowest-energy state in a given charge sector.^{7,24} If we denote the

MPS tensors of the finite chain by $\Gamma^{(1)} \dots \Gamma^{(L)}$,²⁵ then the MPS tensors at the end of the chain [either $\Gamma^{(1)}$ or $\Gamma^{(L)}$] customarily have only two indices, and the state computed by the algorithm is the lowest state having trivial quantum number (i.e. the total charge on all sites is the vacuum). To compute the lowest-energy state in a sector with non-trivial quantum number, one introduces an additional index on either $\Gamma^{(1)}$ or $\Gamma^{(L)}$, and thus on either the system or the environment block. This index behaves similarly to the additional indices introduced in Fig. 5, with the difference that only one such index is introduced, and this index has degeneracy 1.

While the method illustrated in Fig. 5 is capable of finding the translation-invariant state $|\psi'\rangle$ of Eq. (6), it is still less than ideal. State $|\psi'\rangle$ is a ground state with block length 1 in Eq. (4), but requires $D \geq 3$. In contrast, state $|\psi\rangle$ in Fig. 3(i) is a ground state with block length 3, but can be constructed for the lower cost of $D = 1$. More generally, if a ground state may be constructed with period p and bond dimension D , then there also exists a ground state with period 1 and bond dimension pD with construction analogous to $|\psi'\rangle$ in Eq. (6). The approach illustrated in Fig. 5 favours construction of solutions in which the charge configuration of the MPS bond has period 1, and thus to obtain the same level of accuracy as a solution with period p in the charge configuration requires a bond index larger by a factor of p , increasing computational cost by a factor of p^3 . This factor may become quite large.²⁶

The potential weakness of this algorithm is the requirement that the initial boundary charges associated with non-zero entries in λ , which are selected on a finite system of length 2, should also be appropriate on systems of all lengths $L \in \{2n \mid n \in \mathbb{Z}^+\}$. The boundary charges for each iteration are wholly determined by the charges appearing on the MPS bond index of the iteration before, and if early iterations exclude key charges from the bond indices (either because of overly-restrictive bond dimension or because of qualitatively different behaviour in the small- L and large- L regimes) then this may cause failure to converge to the correct ground state. In a worst-case scenario the bond indices may collapse to a single charge sector on some iteration k , and the free boundary charge approach then offers no advantage over standard iDMRG with an appropriate choice of total charge.

An example of failure due to insufficient bond dimension D may be seen by setting $D < 3$ with Hamiltonian (5), prohibiting the construction of the translation-invariant state (6). A more detailed example is provided in Sec. IV.

D. Increasing the inserted block size

Another alternative, suggested in Ref. 8 in the context of protected symmetries associated with particle number, is to modify the size of the initialisation block and subsequent inserted blocks such that the total length of the

growing lattice is always an exact multiple of the filling fraction. This approach requires the filling fraction (or whatever other symmetric property is responsible for the charge periodicity) to be known beforehand, and ideally to be a constrained parameter. It therefore cannot be applied to Hamiltonians such as

$$\hat{H} = -J \sum_i (c_i^\dagger c_{i+1} - c_{i+1}^\dagger c_i) - \mu \sum_i c_i^\dagger c_i, \quad (11)$$

which is the free fermion Hamiltonian written with a chemical potential μ , as the filling fraction for such a Hamiltonian is determined during the optimisation process as a result of competition between the chemical potential and the hopping terms. When compared to the standard approach in which a pair of sites is introduced on each iteration, the additional computational cost of this approach is a factor of $O(d^{p-2})$ where d is the dimension of the Hilbert space of a single site.

E. Infinite DMRG with an accessory charge

1. Description of algorithm

In Sec. III C the local minimum of Hamiltonian (5) corresponding to the empty lattice was avoided by modifying the boundary conditions of the initialisation procedure. Section III D then considered a modification of the iDMRG algorithm which can achieve the same effect, but is possible only under limited conditions. An improvement over both of the above methods may be achieved by modifying the MPS Ansatz as shown in Fig. 6(i).

In this Figure an additional index has been added to the central tensor, $\lambda^{(k)}$, which does not couple to the Hamiltonian and which carries some number of charges $n_q^{(\lambda)}$, each with degeneracy 1. Optimisation proceeds as before, inserting additional sites and computing a new eigentensor v which decomposes into new $A^{(k+1)}$, $B^{(k+1)}$, and $\lambda^{(k+1)}$, only now the λ -tensor has three indices instead of two. Truncation is performed independently on the left and right sides of $\lambda^{(k+1)}$ during the two singular value decompositions shown in Fig. 6(ii). The cost of the algorithm is increased relative to the customary iDMRG algorithm, but only by a factor of $O[n_q^{(\lambda)}]$.

The additional index on the λ -tensor may be given a physical interpretation as an accessory site, decoupled from the Hamiltonian, which is capable of carrying arbitrary charge. As with the extra charges introduced at the left and right of Fig. 5, the presence of this charge relaxes the requirement that the total charges on the sites of the left half-lattice and right half-lattice sum to the vacuum sector.

Since the accessory site is now included on the tensor v , which is variationally optimised during each iteration, this approach avoids the problem of locking in choices of boundary charges during initialisation, which may or




FIG. 6. (i) Infinite MPS with an accessory charge. (ii) Truncation of singular values with an accessory charge. After obtaining the new eigentensor v , an initial decomposition divides v into U , S , and V^\dagger . The first truncation is performed over the singular values in S . Tensor v' is then constructed from S and V^\dagger , then this is decomposed into U' , S' , and V'^\dagger and a second truncation is performed, over the singular values in S' . Tensor $\lambda^{(k+1)}$ is constructed from $U'S'$. (iii) Calculation of charge sector weights.

may not be appropriate in the large- L limit. This algorithm therefore offers greater flexibility than the free boundary charges algorithm of Sec. III C.

2. Comparison with targeting of individual charge sectors

It is useful to compare the accessory site algorithm with the approach discussed in Sec. III C for targeting individual charge sectors in finite DMRG. In both techniques a single additional index is added to the MPS Ansatz, lifting the constraint that the state as a whole must have a quantum number of zero.

When computing the lowest-energy state in a specific charge sector, the additional charge index is customarily placed on either $\Gamma^{(1)}$ or $\Gamma^{(L)}$ and thus is attached to either the system or the environment block. Iteration of the iDMRG algorithm proceeds in the normal fashion, with the total charge constraint imposed by this additional index yielding a state in the desired charge sector. One may target the lowest-energy state across multiple

charge sectors simultaneously by admitting several different charges on this additional index. If the additional index is traced over, then the symmetrised version of this process (where extra indices appear on both the system and the environment) is exactly the procedure described in Sec. III C. If the additional index is not traced over, then this is instead equivalent to performing several simultaneous MPS simulations in parallel, looking for the lowest-energy state in each charge sector (henceforth the *multiple charges method*), with the total bond dimension D being divided either statically or dynamically between these simulations. The fraction of D attributed to a given charge sector is then equivalent to the bond dimension of the effective independent MPS on that charge sector.

The accessory site form of the iDMRG algorithm differs from this approach by putting the extra index on the central tensor λ , rather than on the system or environment. This simple modification has a number of very important effects.

1. In the multiple charges method, superselection requires that for an MPS bond dimension D , the DMRG algorithm must divide this bond dimension between the multiple different charge sectors being independently optimised. In contrast, accessory site iDMRG typically selects a single charge sector yielding the lowest-energy state and allocates the entire MPS bond dimension to that state.²⁷ For conventional iDMRG this would result in the MPS becoming locked into a single charge sector. This does not happen for accessory site iDMRG, however, because of the second major difference between the multiple charges method and the accessory site method:
2. In the multiple charges method, each of the parallel MPS simulations evolves independently. Thus for a simulation with total charge q , the system and environment blocks at step k correspond to the optimal system and environment superblocks *for total charge q* at step $k - 1$. In accessory site iDMRG, the system and environment blocks at step k correspond to the optimal system and environment superblocks *across all charges* at step $k - 1$, even if the optimal state at step k has total charge q and the optimal state at step $k - 1$ has total charge $q' \neq q$.

As a result, accessory site iDMRG is capable both of converging more rapidly to a lower-energy configuration than the multiple charges method (as is clearly seen in the example in Sec. IV), and of yielding a more faithful approximation to the true ground state (corresponding to a higher effective bond dimension D in the optimal charge sector at each iteration).

3. Interpretation of results

To interpret the results obtained using this Ansatz, it is necessary to compute the relative weights of the different charge sectors on the accessory site for several iterations. These are given by the diagonal elements of the matrix

$$M_{q'}^q = \lambda_{\beta}^{\alpha q} \lambda_{\alpha q'}^{\dagger \beta} \quad (12)$$

[Fig. 6(iii)]. The behaviour of these values then falls into one of the following four categories:

1. *Single charge sector*—Only one charge sector has a weight differing from zero by more than can be attributed to machine precision rounding errors, and this sector is consistent from iteration to iteration. The lowest-energy state identified after k iterations takes the form of Eq. (4), with possible periodicity as per Refs. 14–15 and $\lambda^{(k)}$ as per Fig. 6(i). The tensor $\lambda^{(k)}$ carries the specified charge q on the accessory site. If this charge is not the vacuum charge, consider the meaning of this accessory charge in the context of your model. If you wish to find the next-lowest-energy charge sector, repeat the simulation but exclude this charge from the new index of the λ -tensor. Note that if the only non-trivial sector has the vacuum charge on the accessory site, this site can be projected down to the vacuum sector and then deleted to recover a standard MPS without an accessory site.
2. *Superposition of charge sectors*—When more than one charge sector has a weight different to zero, this indicates the existence of multiple degenerate ground states with different charges on the accessory site. As the ground states are degenerate, the relative occupancies of the charge sectors may vary from iteration to iteration even while the MPS is otherwise converged, unless fixed by a choice of gauge. If wishing to restrict to a single charge value on the accessory site, either project onto this sector and rescale the relevant singular values, or exclude all but one of these sectors and repeat the simulation. The absence of a superposition does not confirm the absence of a degeneracy.
3. *Cyclic charge sectors*—Only one charge sector has a weight different to zero on any given iteration, but this charge changes cyclically from iteration to iteration with a period p_q . To identify a ground state with charge q on the accessory site, halt the simulation when only sector q has non-trivial weight. The Ansatz for the ground state is now given by Eq. (4) with a block length p which is an integer multiple of p_q , being the lowest common multiple of p_q and any periodicity of the sort described in Refs. 14–15. Again note that if the accessory charge is trivial, the accessory site can be projected then deleted.

4. Cyclic superpositions of charge sectors—A combination of outcomes 2 and 3.

With limited exceptions (see e.g. Sec. III F), the desired outcome will be an Ansatz for a state with vacuum charge on the accessory site. However, by permitting nontrivial charges on the accessory site during the optimisation process one may obtain states such as that of Fig. 3(i) which cannot be reached by the usual symmetric iDMRG algorithm. For example, given Hamiltonian (5) and \mathbb{Z}_3 symmetry one obtains a result with cyclic charge sectors (outcome 3). Halting when the non-zero weight is in the vacuum sector, deleting the accessory site, and applying Eq. (4) with block length $p = p_q$ gives state $|\psi\rangle$ as shown in Fig. 3(i), with MPS bond dimension $D = 1$.

Note—A procedure for generating an initial estimate for vector v in conventional iDMRG is given in Ref. 28. This procedure converges to the same fixed point as the Product Wave-Function Renormalisation Group technique (PWFRG) for TEBD,^{29–32} and may be readily generalised to the accessory site algorithm (see Appendix A). However, it is unable to allow for situations where the charge on the accessory site is cyclic. When this is the case, the overlap of the initial guess with the true lowest-energy eigenvector may be sufficiently small as to result in failure of the numerical algorithm (e.g. Lanczos) used to find v . For this reason, the use of this estimation procedure is not recommended with the current algorithm unless the charge on the accessory site is fixed by the specification of the system to be studied.

F. Topologically ordered systems

Occasionally, situations may be encountered where it is appropriate to apply an explicit physical interpretation to the accessory site. An example of this is in the study of topologically ordered systems with a boundary. Models of this sort include Kitaev’s toric code Hamiltonian³³ on a cylinder, critical models whose infrared limit corresponds to a topological quantum field theory (TQFT),^{34,35} and models whose particle excitations explicitly correspond to the charges of a TQFT or unitary braided tensor category.^{19,36–38}

These models admit dual descriptions either in terms of the microscopic degrees of freedom, such as the spins of Kitaev’s toric code, or in terms of the emergent topological order. For the toric code, stars and plaquettes which are not in a lowest-energy configuration may be associated with e and m charges respectively, and may be created pairwise (two e charges or two m charges) by means of single spin flips in the ground state. Their product is the fermionic excitation em , and these three charges, together with the vacuum charge \mathbb{I} , exhibit the fusion rules and exchange statistics of the symmetry group $D(\mathbb{Z}_2)$. More generally, the charges of an emergent topological order on a 2D manifold may exhibit fusion and braiding statistics derived from any unitary braided tensor category (though this is restricted to unitary *modular* braided

tensor categories if the manifold supports a non-trivial cycle). These generalised topologically ordered quasi-particle excitations are termed *anyons*,^{16–18} and are of particular interest due to their potential use in quantum computation³⁹ and their proposed role in explaining the plateau states of the fractional quantum Hall effect.^{40–45} This interest is likely to be further enhanced by the recently reported detection of Ising anyons at the ends of iron nanowires.⁴⁶ Tensor networks provide one of the most promising approaches for the study of large-scale anyonic systems,^{36,37} and therefore the application of accessory site iDMRG to anyonic systems is of particular importance.

When simulating any topologically ordered system using a tensor network algorithm, it is possible to protect the symmetries associated with the topological charges of the model. In the enforced absence of symmetry-breaking operations the topological charges behave as good quantum numbers, and may be exploited in tensor network algorithms for computational gain (and sometimes, in the absence of an explicit microscopic model, as a matter of necessity). The application of this approach to the finite DMRG algorithm is presented in Ref. 47, and builds on the earlier work of Ref. 36 to develop a general toolbox for the construction and manipulation of anyonic tensors. This approach to DMRG for topologically ordered systems (including anyons) may be viewed as a generalisation of the exploitation of non-Abelian symmetries seen in earlier DMRG simulations.⁹

While the toric code is most commonly studied (surprise, surprise) on the torus, other anyonic lattice systems are frequently studied on manifolds with a boundary, most commonly (and sometimes by implicit assumption) on the disc. However, any manifold boundary is topologically equivalent to one or more punctures, and thus for topologically ordered systems any manifold with a boundary admits a (total) boundary charge. Much as enforcing a total spin sector restricts the admissible states in conventional DMRG, the choice of this charge may affect the family of states admissible on the lattice, even if no explicit Hamiltonian-mediated coupling between the lattice and the boundary exists.

For finite systems, as discussed in Ref. 47, this boundary charge may be included as an additional site at one end of the chain, and this approach is directly analogous to the selection of a specific spin sector in conventional DMRG. The naïve generalisation of this approach to infinite systems is to include this charge on one side or the other of the initial state, as per Fig. 5, though we have already seen that this is suboptimal. The boundary charge may then be fixed by only admitting a single choice of charge on this boundary index, or multiple possible boundary charges may be admitted.

Once again, this treatment is improved upon by the approach described in Sec. III E where the boundary charge is expressed on the λ -tensor. As previously discussed, this results in a decrease in computational cost when the ground state may be expressed in block-periodic form (4),




FIG. 7. Mean filling fraction for a system of free hard-core \mathbb{Z}_3 anyons on a ladder (13), after 300 iterations, as a function of μ/t . Results are shown for conventional iDMRG (+), iDMRG with free boundary charges (\times), and accessory site iDMRG (\circ), all with MPS bond dimension $D = 100$. Both iDMRG with free boundary charges and accessory site iDMRG correctly identify full occupation of the lattice at large values of μ/t , but small discrepancies between these techniques are still observed at intermediate fillings.

and gives the ability to explicitly determine the boundary charge of the ground state by calculating weights as per Fig. 6(iii), and to restrict it by constraining the admissible charges on the accessory site. Finally, this approach delivers a further benefit when the boundary charge admits a real, physical interpretation: If the optimal boundary charge differs in the low- L and high- L regimes, then the inclusion of the accessory site in every update step facilitates the transition between these regimes, whereas if the boundary charge is included on the system or environment block as per Fig. 5, the qualitative behaviour of the low- L regime may persist as charge boundary conditions are locked in by unfortunate choices of initial tensors $A^{(1)}$ and $B^{(1)}$.

IV. FURTHER EXAMPLE

As a more realistic example of the application of accessory site iDMRG, consider a system of hard-core \mathbb{Z}_3 anyons with variable filling on a two-rung ladder. The only interactions are a term permitting anyons to hop to neighbouring sites if these are vacant, and a chemical potential. Braiding arises when two anyons use vacant spaces to hop around one another. Schematically, the Hamiltonian may be written

$$\hat{H} = -t \sum_{\langle i,j \rangle} (\text{hopping terms}) - \mu \sum_i c_i^\dagger c_i \quad (13)$$




FIG. 8. Mean filling fraction $\langle n \rangle$ as a function of the number of iterations performed, for conventional iDMRG (+), iDMRG with free boundary charges (\times), and accessory site iDMRG (\circ), all at $\mu/t = 1.5$ with MPS bond dimension $D = 100$.

where $\langle i, j \rangle$ represents all pairs of neighbouring sites. Physical indices of the MPS may be in charge sector 0 or 1 corresponding to absence or presence of an anyon respectively, while MPS bond indices carry the cumulative anyonic charge, summing from left to right, which may be 0, 1, or 2. All tensors are symmetric, with nonzero entries only for combinations of charge labels whose oriented sum is zero. Figure 7 shows filling fraction as a function of μ/t after three hundred iterations. Conventional iDMRG with exploitation of \mathbb{Z}_3 symmetry is immediately seen to be unreliable at high filling fractions, for the reasons discussed in Sec. III A. Infinite DMRG with boundary charges and accessory site iDMRG both fare better at high filling fractions, correctly capturing the $\langle n \rangle = 1$ insulating plateau, but show disagreement at intermediate fillings.

To understand this discrepancy, let μ/t be fixed at 1.5 and consider how the value of $\langle n \rangle$ evolves as a function of the number of DMRG iterations performed, as shown in Fig. 8. At $\mu/t = 1.5$ the filling fraction is sufficiently low that even the standard symmetric iDMRG algorithm does not suffer a catastrophic failure, yielding a superficially reasonable value for $\langle n \rangle$. Nevertheless, convergence of the standard algorithm is slow, exhibiting an extremely long periodicity of the sort described in Refs. 14–15. The oscillations in $\langle n \rangle$ are still large after 300 iterations, and they exhibit noticeable variability from one cycle to the next, suggesting that the simulation is relatively poorly converged.

In contrast the periodicity for accessory site iDMRG is much shorter and the rate of convergence is faster, giving more accurate results after significantly fewer iterations.

Comparison of the mean occupation number with a reference value obtained using a non-symmetric implementation of iDMRG is revealing: Protection of symmetry has forced the standard (symmetric) iDMRG algorithm to converge towards a much longer-period orbit than either the non-symmetric iDMRG algorithm or the accessory charge iDMRG algorithm (Fig. 8), and the mean value of the filling fraction over this orbit differs from that obtained using the non-symmetric reference implementation (Table I). This larger orbit must therefore be considered a symmetry-protected local minimum, while the non-symmetric iDMRG and accessory site iDMRG converge to orbits which more accurately approximate the true ground state.

Interestingly, at $\mu/t = 1.5$ the free boundary charge algorithm is seen to settle into a similar longer-period orbit to the standard symmetric algorithm. For the free boundary charge algorithm this represents a relatively benign form of failure due to insufficiently large bond dimension and/or overconstrained boundary conditions at small chain lengths, as described in Sec. III C. Rather than recovering a translation-invariant representation of the ground state as per Eq. (6), the algorithm has broken translation symmetry in pursuit of lower energies. The benefit associated with the additional boundary charges has been lost, and the algorithm has ended up in a similar symmetry-protected local minimum to the standard symmetric iDMRG. Thus, while the free boundary charge algorithm is equally as effective as the accessory charge algorithm at sufficiently high and sufficiently low values of μ/t , it is less reliable over an intermediate regime.

The generalisation of these observations to arbitrary Hamiltonians is as follows:

- Standard symmetric iDMRG: Avoids symmetry-protected local minima if the maximum achievable value of $\hat{\rho}^{(P)}$ in Eq. (8) is sufficiently small, as charge boundary conditions during initialisation of the MPS are not of critical importance.
- Free boundary charge iDMRG: Avoids symmetry-protected local minima if the maximum achievable value of $\hat{\rho}^{(P)}$ is either sufficiently small, or sufficiently large and accompanied by a sufficiently large MPS bond dimension D . These conditions

promote the retention of all significant charge sectors on the MPS bond. When the free boundary charge algorithm does become stuck in a local minimum, the effects are likely to be less severe as effective convergence may still be attained on the portions of the MPS tensors not participating in the intertwinors, when working in the intermediate $\hat{\rho}^{(P)}$ regime.

- Accessory site iDMRG: Efficiently avoids symmetry-protected local minima for any maximum achievable value of $\hat{\rho}^{(P)}$ and any MPS bond dimension D .

V. IMPLICATIONS FOR iTEBD

As explained in Sec. III of Ref. 28, there exists a deep connection between the iDMRG algorithm and the infinite Time-Evolving Block Decimation algorithm (iTEBD) with imaginary time evolution. Although these algorithms differ substantially in their usual presentation, they are in fact formally equivalent up to a choice of eigensolver used to find the ground state, with iDMRG employing the Lanczos method where iTEBD uses the power method, repeatedly applying an infinitesimal imaginary time evolution operator to a tensor pair then restoring normalisation of the wavefunction.

As a consequence, precisely the same local minima exist in symmetric iTEBD as have been seen here for iDMRG. The iTEBD algorithm, however, is much more robust against becoming stuck in a significant symmetry-protected local minimum. This is because the most common initialisation procedure for iTEBD is the selection of a random pair of initial tensors A and B . The Ansatz for the infinite chain then comprises the pair AB repeated an infinite number of times. The boundary conditions of this initialisation are equivalent to those given in Sec. III C for iDMRG, with all valid charges potentially being represented on the boundary. For sufficiently large MPS bond dimension D , the iTEBD algorithm will therefore converge towards a state having a form analogous to Eq. (6).

Nevertheless, as seen in Sec. IV, it is still possible for iDMRG with free boundary charges—and therefore also for symmetric iTEBD—to become stuck in a symmetry-protected local minimum. However, this risk is only significant for intermediate values of $\hat{\rho}^{(P)}$, implying that only a part of each MPS tensor in the translation-invariant ground state exhibits an intertwinor structure. As seen in the example of Table I, relatively good convergence to the true ground state is therefore possible, though the end precision is still less than that which can be obtained using accessory site iDMRG.

In summary, the standard iTEBD algorithm is therefore largely protected from the most damaging symmetry-protected local minima, but potentially performs suboptimally both in speed of convergence and in the precision obtainable for a given MPS bond dimen-

TABLE I. Mean filling fraction for Hamiltonian (13) with $\mu/t = 1.5$ after 300 iterations, for different variants of the iDMRG algorithm with MPS bond dimension $D = 100$. The average $\langle n \rangle$ is taken over one orbit within the space of MPS states.

Algorithm	$\langle n \rangle$
Reference non-symmetric iDMRG	0.73335
Standard symmetric iDMRG	0.73602
Free boundary charges iDMRG	0.73636
Accessory site iDMRG	0.73336

sion D . It is, however, unclear whether a counterpart to the accessory site iDMRG algorithm exists for iTEBD, as a necessary prerequisite for accessory site iDMRG is a numerical eigensolver which can switch between charge sectors from one iteration to the next. If this capability is suppressed, for example by using the initialisation vector estimation process of Appendix A, then the accessory site algorithm is no better than standard iDMRG. At this time it is unclear how this ability to switch charge sectors might be incorporated into the iTEBD algorithm, and thus it is uncertain whether the iTEBD algorithm can be modified to obtain the enhanced performance seen in accessory site iDMRG.

VI. CONCLUSION

Under certain circumstances, described in Sec. III B, exploitation of symmetries in infinite DMRG may cause failure of simulations to converge to the correct ground state, becoming trapped instead in *symmetry-protected local minima*. This paper has examined three modified forms of the infinite DMRG algorithm which permit reliable exploitation of symmetries of the Hamiltonian at high filling fraction. Of these, one (discussed in Sec. III D and Ref. 8) requires *a priori* knowledge of the repeated block length of the converged symmetric MPS and so is limited in its application. Another method (Sec. III C) is a generalisation of techniques used in finite DMRG to identify states having non-zero quantum number. This method performs reasonably well for an example realistic Hamiltonian, but may still encounter problems if the MPS bond dimension is insufficiently large, or if the boundary charges which are appropriate at initialisation are not well-suited to longer chains. The third method (Sec. III E) attaches an accessory site to the region undergoing variational update. This method successfully reproduces expected behaviour in all regimes, robustly re-evaluates the global boundary conditions on each iteration, offers improved rates of numerical convergence, and frequently requires a much smaller MPS bond dimension than the preceding method to obtain the same level of precision. When exploiting symmetries in iDMRG, accessory site iDMRG is consequently the method of choice.

From the relationship between iDMRG and iTEBD described in Ref. 28, the boundary conditions in symmetric iTEBD are seen to be equivalent to those introduced in Sec. III C. A more comprehensive discussion may be found in Secs. IV and V, but in brief it follows that while iTEBD is relatively robust against symmetry-protected local minima, it may deliver results with precisions inferior to those of accessory site iDMRG for equivalent MPS bond dimension.

The accessory site algorithm also permits the study of systems where the extra site admits a physical interpretation, such as topologically ordered (e.g. anyonic) systems on a manifold with a boundary, for which it represents the total boundary charge. In such systems the acces-




FIG. 9. (i) Decomposition of λ into a unitary matrix U , a diagonal matrix S which carries the singular values, and an isometry T which carries the accessory site. (ii) Rolling tensor λ left through tensor A . The decomposition is chosen such that B^L is right-unitary as per Fig. 1(i). (iii) The left inverse of λ^0 . (iv) Construction of λ^T using the left inverse of λ^0 .

sory charge may be left free (with its admissible value(s) in the ground state(s) being an output of the algorithm), or partially or fully constrained. The standard infinite DMRG algorithm (without an accessory site or boundary charges) is recovered if this site is constrained to hold the trivial charge.

Appendix A: Estimation of iDMRG update vector

This Appendix describes how the prescription given in Ref. 28 for estimating the eigenvector v may be generalised when v has an additional index corresponding to an accessory site. Its use is not recommended unless the charge on this accessory site is fixed by the system under study, as it is unable to accommodate cycling of the charge on the accessory site. Use of this prescription

with the accessory site algorithm has been observed to cause simulations to converge to a state with fixed accessory charge when a lower-energy series of configurations with cyclic accessory charge (and the same MPS bond dimension) is known to exist.

To generalise the prescription of Ref. 28, first let λ be decomposed as shown in Fig. 9(i). As a choice of gauge the unitary matrix may always be absorbed into the A -tensor to the immediate left, and thus λ comprises a diagonal matrix S containing singular values, and an isometry T (a projector which obeys unitarity on the smaller Hilbert space) which carries the accessory site.

As in Ref. 28, the process at the heart of the update procedure is “rolling” left and right, in which the location of the λ -tensor moves one bond to the left or the right on the MPS chain. To perform the process of “rolling left”, for example, the λ -tensor and the immediately-adjacent A -tensor are fused, then separated again by means of singular value decomposition or QR decomposition as shown

in Fig. 9(ii). (Details such as the sequence of fusing and splitting of indices, while necessary, are routine and so have not been explicitly specified.) The new λ -tensor, λ^L , is again gauged so that $U = \mathbb{I}$ in Fig. 9(i). Similarly, λ^R is generated by rolling right. Finally, λ^0 denotes the optimal λ -tensor computed on the iteration before last.

It then suffices to construct either the left inverse or right inverse of λ^0 , with the left inverse being shown in Fig. 9(iii), and then constructing the trial vector

$$\lambda^T = \lambda^R (\lambda^0)^{-1} \lambda^L \quad (\text{A1})$$

where the charge index of $(\lambda^0)^{-1}$ is summed with that of λ^R if using the right inverse, or with that of λ^L if using the left inverse [Fig. 9(iv)]. The trial vector is then assembled in the usual way from λ^T and the A - and B -tensors generated when rolling right and left (and incorporating any unitary matrices generated when gauging λ^L and λ^R).

* robert.pfeifer@mq.edu.au

¹ S. R. White, *Phys. Rev. Lett.* **69**, 2863 (1992).

² S. R. White, *Phys. Rev. B* **48**, 10345 (1993).

³ U. Schollwöck, *Rev. Mod. Phys.* **77**, 259 (2005).

⁴ U. Schollwöck, *Ann. Phys.* **326**, 96 (2011).

⁵ S. Östlund and S. Rommer, *Phys. Rev. Lett.* **75**, 3537 (1995).

⁶ F. Verstraete, D. Porras, and J. I. Cirac, *Phys. Rev. Lett.* **93**, 227205 (2004).

⁷ I. P. McCulloch, *J. Stat. Mech.: Theory and Exp.* **2007**, P10014 (2007).

⁸ I. P. McCulloch and M. Gulácsi, *Australian J. Phys.* **53**, 597 (2000).

⁹ I. P. McCulloch and M. Gulácsi, *Europhys. Lett.* **57**, 852 (2002).

¹⁰ S. Singh, R. N. C. Pfeifer, and G. Vidal, *Phys. Rev. A* **82**, 050301 (2010).

¹¹ S. Singh, R. N. C. Pfeifer, and G. Vidal, *Phys. Rev. B* **83**, 115125 (2011).

¹² R. N. C. Pfeifer, *Simulation of Anyons Using Symmetric Tensor Network Algorithms*, Ph.D. thesis, The University of Queensland (2011), arXiv:1202.1522v2 [cond-mat.str-el].

¹³ S. Singh and G. Vidal, *Phys. Rev. B* **86**, 195114 (2012).

¹⁴ S. Caprara and A. Rosengren, *Europhys. Lett.* **39**, 55 (1997).

¹⁵ S. Caprara and A. Rosengren, *Nuc. Phys. B* **493**, 640 (1997).

¹⁶ A. Kitaev, *Ann. Phys.* **321**, 2 (2006).

¹⁷ P. H. Bonderson, *Non-Abelian Anyons and Interferometry*, Ph.D. thesis, California Institute of Technology (2007).

¹⁸ P. Bonderson, K. Shtengel, and J. Slingerland, *Ann. Phys.* **323**, 2709 (2008).

¹⁹ R. N. C. Pfeifer, O. Buerschaper, S. Trebst, A. W. W. Ludwig, M. Troyer, and G. Vidal, *Phys. Rev. B* **86**, 155111 (2012).

²⁰ For $\varepsilon = 0$ and infinite MPS bond dimension D it is in theory possible to construct a Hamiltonian for which p is

infinite; an example would be a system whose filling fraction tends towards an irrational number as $k \rightarrow \infty$, with particle number a protected symmetry. By definition such a system does not converge for any finite k . In practice, with finite D and ε and a deterministic computing device, the state space of the calculation is finite and any trajectory through this space will eventually repeat for some p , which may or may not be small enough to be recognised in practice.

²¹ R. N. C. Pfeifer, arXiv:1505.06928 [cond-mat.str-el] (2015).

²² The notion of locality on the space of matrix product states is not, in general, rigorously defined. A state $|\psi'\rangle$ may be thought of as local with respect to a state $|\psi\rangle$, a Hamiltonian \hat{H} , and an algorithm \mathcal{A} if a single iteration of algorithm \mathcal{A} with Hamiltonian \hat{H} is capable of taking $|\psi\rangle$ into $|\psi'\rangle$.

²³ If \hat{H} is invariant under translation by p sites, a translation-invariant representation of \hat{H} may be constructed by fusing p sites of dimension d into one effective site of dimension d^p .

²⁴ I. McCulloch, M. Gulácsi, S. Caprara, A. Jazavaou, and A. Rosengren, *J. Low Temp. Phys.* **117**, 323 (1999).

²⁵ G. Vidal, *Phys. Rev. Lett.* **93**, 040502 (2004).

²⁶ Note that the algorithm as a whole may still display periodicity—it is only periodicity of charge configurations which is suppressed and which causes increased computational cost. For example, in Refs. 14–15 conservation of particle number is not exploited for computational gain. The results are asymptotically p -periodic with $p > 1$, but all indices carry only the trivial charge.

²⁷ A superposition of charge sectors may be selected if both yield equally low energy states, and D is redundantly large.

²⁸ I. P. McCulloch, arXiv:0804.2509v1 [cond-mat.str-el] (2008).

²⁹ T. Nishino and K. Okunishi, *J. Phys. Soc. Jpn.* **64**, 4084 (1995).

³⁰ Y. Hieida, K. Okunishi, and Y. Akutsu, *Phys. Lett. A* **233**, 464 (1997).

³¹ K. Okunishi, Y. Hieida, and Y. Akutsu,

- [Phys. Rev. B](#) **59**, 6806 (1999).
- ³² K. Ueda, T. Nishino, K. Okunishi, Y. Hieida, R. Derian, and A. Gendiar, [J. Phys. Soc. Jpn.](#) **75**, 014003 (2006).
- ³³ A. Y. Kitaev, [Ann. Phys.](#) **303**, 2 (2003).
- ³⁴ P. Di Francesco, P. Mathieu, and D. Sénéchal, *Conformal Field Theory* (Springer, New York, 1997).
- ³⁵ R. N. C. Pfeifer, G. Evenbly, and G. Vidal, [Phys. Rev. A](#) **79**, 040301 (2009).
- ³⁶ R. N. C. Pfeifer, P. Corboz, O. Buerschaper, M. Aguado, M. Troyer, and G. Vidal, [Phys. Rev. B](#) **82**, 115126 (2010).
- ³⁷ R. König and E. Bilgin, [Phys. Rev. B](#) **82**, 125118 (2010).
- ³⁸ R. N. C. Pfeifer, [Phys. Rev. B](#) **85**, 245126 (2012).
- ³⁹ C. Nayak, S. H. Simon, A. Stern, M. Freedman, and S. Das Sarma, [Rev. Mod. Phys.](#) **80**, 1083 (2008).
- ⁴⁰ N. Read, [Phys. Rev. Lett.](#) **65**, 1502 (1990).
- ⁴¹ N. Read and G. Moore, [Prog. Theo. Phys. Supp.](#) **107**, 157 (1992).
- ⁴² N. Read and E. Rezayi, [Phys. Rev. B](#) **59**, 8084 (1999).
- ⁴³ J. S. Xia, W. Pan, C. L. Vicente, E. D. Adams, N. S. Sullivan, H. L. Stormer, D. C. Tsui, L. N. Pfeiffer, K. W. Baldwin, and K. W. West, [Phys. Rev. Lett.](#) **93**, 176809 (2004).
- ⁴⁴ W. Pan, J. S. Xia, H. L. Stormer, D. C. Tsui, C. Vicente, E. D. Adams, N. S. Sullivan, L. N. Pfeiffer, K. W. Baldwin, and K. W. West, [Phys. Rev. B](#) **77**, 075307 (2008).
- ⁴⁵ A. Kumar, G. A. Csáthy, M. J. Manfra, L. N. Pfeiffer, and K. W. West, [Phys. Rev. Lett.](#) **105**, 246808 (2010).
- ⁴⁶ S. Nadj-Perge, I. K. Drozdov, J. Li, H. Chen, S. Jeon, J. Seo, A. H. MacDonald, B. A. Bernevig, and A. Yazdani, [Science](#) **346**, 602 (2014).
- ⁴⁷ R. N. C. Pfeifer and S. Singh, [arXiv:1505.00100v1 \[cond-mat.str-el\]](#) (2015).